arXiv:1610.06996v1 [cs.LO] 22 Oct 2016

Automation of separation logic using auto2

Bohua Zhan

Massachusetts Institute of Technology

Abstract. We present a new system of automation for separation logic in the interactive theorem prover Isabelle. The system is based on the recently developed auto2 prover, and follows a natural, saturation-based approach to reasoning about imperative programs. In addition to standard examples on linked lists and binary search trees, we apply the automation to red-black trees and indexed priority queues, showing that it provides a high degree of automation even on the more complicated data structures.

1 Introduction

Separation logic [20] is an extension of Hoare logic designed to facilitate reasoning about heap-based programs. It allows one to reason about such programs locally: the Hoare triple of each command is only applied to the part of the heap that the command uses, while the other parts remain unchanged.

There has been extensive work on automation of separation logic, in both automatic and interactive theorem proving settings. Our focus is on automation in the interactive setting. In this setting, there are usually no constraints on the logic used, and arbitrarily complex algorithms can be verified in theory. On the other hand, human guidance, often requiring substantial expertise, is needed for the more complex verifications. The focus of automation in this area is to reduce the amount of effort and expertise needed for the verification tasks that occur in practice.

In this paper, we describe a new system of automation for separation logic, based on the recently introduced auto2 prover [22], within the proof assistant Isabelle. By writing the automation as an extension to auto2, we naturally use a saturation-based, best-first search mechanism, in contrast to the depth-first search characteristic of automations written using the tactics framework. Moreover, we aim to emulate a style of proof that is very close to human reasoning: by working directly with the heaps and the data structures on them, and accumulate facts about these objects.

For program logic, we use the Imperative-HOL framework [3] in the Isabelle library. In this framework, programs are represented using Haskell-style monads. By natively supporting arrays and storing records containing several fields at a single address, it also removes any need for pointer arithmetic. Nevertheless, the focus of our work is orthogonal to such simplifications, and the same approach should be applicable to program logics with more realistic memory models and

C-style semantics. A separation logic for Imperative-HOL, together with an automation framework based on tactics, is developed by Lammich and Meis [11]. We will use the construction of separation logic, but not the automation, in their work.

Auto2 is a general-purpose theorem prover for classical higher-order logic implemented in Isabelle. It makes progress by repeatedly calling a set of *proof steps*, which are reasoning rules written as ML functions. By extending this set of proof steps (and making other customizations), one can implement potentially complicated, domain-specific heuristics in auto2. Another purpose of this paper is to demonstrate this extensibility, using separation logic as an example. We review the relevant concepts from Imperative-HOL, its separation logic, and from auto2 in Section 2.

Given a Hoare triple to prove, our method creates variables for each heap encountered during the execution of the procedure, and for each data structure on those heaps. It then accumulates facts about these variables, using the pre-condition as the assumption, the post-condition as the goal, and previously proved Hoare triples as reasoning rules. The search mechanism allows multiple Hoare triples for the same procedure, so that complex specifications of a procedure can be verified incrementally. We give a detailed description of the system in Section 3.

We apply our framework to verify the correctness of several imperative data structures. Besides the usual examples of linked lists and binary search trees, we verify two advanced data structures: red-black trees and indexed priority queues. In all these examples, the automation is able to take care of nearly all reasoning concerning separation logic, so that most of the work by the user is actually concerned with abstract data structures. The examples are described in Section 4.

In Section 5, we compare our approach to other approaches for automating separation logic. Finally, we conclude in Section 6 and discuss some future directions of research.

2 Background

In this section, we review concepts from Imperative-HOL, its separation logic, and auto2 that are needed for the rest of the paper.

2.1 Imperative-HOL and separation logic

We begin by reviewing some basic concepts from Imperative-HOL and its separation logic. See [3,10,11] for details.

Heaps and programs In the Imperative-HOL framework, procedures are represented using Haskell-style monads. They operate on a heap (type `heap`) consisting of a finite mapping `ref` from addresses (natural numbers) to natural numbers, and a finite mapping `arrays` from addresses to lists of natural numbers (in order

to support arrays). The heap also contains a counter lim that specifies the next address at which both the ref and arrays mappings are empty. New values and arrays are allocated to position lim, so that any address less than lim contains exactly one value or array. Values of any type α can be stored in the heap by specifying an injection from type α to the natural numbers. This means records with multiple fields, such as nodes of a search tree, can be stored in a single address. Along with native support for arrays, this eliminates any need for pointer arithmetic. Note also that there is no separate stack in a monadic framework.

The type of a procedure returning a value of type α is given by

$$\texttt{datatype}\ \alpha\ \mathsf{Heap} = \mathrm{Heap}\ \mathsf{heap} \Rightarrow (\alpha \times \mathsf{heap})\ \mathsf{option}.$$

The procedure takes as input a heap h, and outputs either None for failure, or Some (r, h'), where r is the return value and h' is the new heap. The bind function for sequencing two procedures has the usual type α Heap $\to (\alpha \to \beta$ Heap$) \to \beta$ Heap. In this paper, we will write $r \leftarrow c1; c2$ for bind $c1$ $(\lambda r.\, c2)$.

Imperative-HOL does not have a native support for while loops. Instead, recursion is used throughout, with properties of recursive procedures proved by induction. We will follow the same approach in our examples.

Partial heaps A *partial heap* is an element of type pheap = heap × nat set. The partial heap (h, as) represents the part of heap h given by the set of addresses as. A partial heap is valid if all elements of as are less than lim h.

The crucial concept of successful execution on a partial heap is defined as follows: an execution of procedure c on a partial heap (h, as) is successful if the execution of c on h returns Some (r, h'), the heaps h and h' agree on all addresses within the limit of h and outside the address set as, and the limit of h' is at least the limit of h. That is, the procedure c does not modify any address outside as, except possibly allocating new memory.

Assertions An assertion (type assn) is a mapping pheap ⇒ bool, with the condition that it is true only on valid partial heaps, and it does not depend on values of the heap outside the address set. The notation $(h, as) \vDash P$ stands for "the assertion P holds on the partial heap (h, as)".

The primitive assertions are defined and then proved to satisfy the required condition. They include:

- true: holds for all valid partial heaps.
- emp: the partial heap is empty.
- $\uparrow (b)$: the partial heap is empty and b (a boolean value) holds.
- $p \mapsto_r a$: the partial heap contains a single address pointing to value a.
- $p \mapsto_a xs$: the partial heap contains a single address pointing to list xs.

The *separating conjunction* on two assertions is defined as follows:

$$P * Q = \lambda(h, as).\, \exists u\ v.\, u \cup v = as \wedge u \cap v = \emptyset \wedge (h, u) \vDash P \wedge (h, v) \vDash Q.$$

This gives an associative and commutative (AC) operator with unit emp. The existence quantification on assertions is defined as:

$$\exists_A x.\, P(x) = \lambda(h, as).\, \exists x.\, (h, as) \models P(x).$$

Assertions of the form $\uparrow(b)$ are called *pure* assertions. In [11], conjunction, disjunction, and the magic wand operator on assertions are also defined, but we will not use them here.

Hoare triples A Hoare triple is a predicate of type

$$\mathsf{assn} \Rightarrow \alpha\ \mathsf{Heap} \Rightarrow (\alpha \Rightarrow \mathsf{assn}) \Rightarrow \mathsf{bool}$$

defined as follows: $\langle P\rangle\, c\, \langle Q\rangle$ holds if for any partial heap (h, as) satisfying P, the execution of c on the partial heap (h, as) is successful with new heap h' and return value r, and the new partial heap (h', as') satisfies $Q(r)$, where $as' = as \cup \{x.\, \mathsf{lim}\ h \leq x < \mathsf{lim}\ h'\}$.

With these definitions, it is possible to prove the *frame rule*:

$$\langle P\rangle\, c\, \langle Q\rangle \implies \langle P * R\rangle\, c\, \langle \lambda x.\, Q(x) * R\rangle.$$

Example: ordered insertion on linked lists We will use linked lists and ordered insertion as the running example of this paper. The example shown here is a slight modification of our case study on linked lists, which in turn follows the corresponding example in [11]. The assertion for linked lists is defined inductively as follows:

```
os_list [] p = ↑(p = None)
os_list (x :: xs) (Some p) = ∃Aq. p ↦r Node x q ∗ os_list xs q
os_list (x :: xs) None = false
```

The os_prepend procedure, given below, simply inserts an element to the front of a linked list:

```
os_prepend a n = do { p ← ref (Node a n); return (Some p) }
```

The os_insert procedure inserts an element to its proper place, assuming the input list is sorted. Note this is a recursive definition.

```
os_insert x b = (case b of
    None ⇒ os_prepend x None
  | Some p ⇒ do {
      v ← !p;
      (if x ≤ val v then os_prepend x b
       else do {
         q ← os_insert x (nxt v);
         p := Node (val v) q;
         return (Some p) }) })
```

Two important properties of os_insert are as follows: First, it acts as expected on the set of elements:

```
<os_list xs b>
os_insert x b
<λr. ∃A xs'. os_list xs' r ∗ ↑(set xs' = {x} ∪ set xs)>
```

Next, it preserves sortedness:

```
<os_list xs b ∗ ↑(sorted xs)>
os_insert x b
<λr. ∃A xs'. os_list xs' r ∗ ↑(sorted xs')>
```

Note we stated these properties as separate Hoare triples. Of course, in this case it is easy to give a functional specification of os_insert. However, we avoid this approach for illustrative purposes.

2.2 The auto2 prover

Auto2, introduced recently in [22], is a general-purpose theorem prover for classical higher-order logic in Isabelle. It is designed to be used in an interactive setting: the user can easily add new heuristics and custom procedures to the prover, as well as provide hints in the middle of a proof. At the same time, it uses a robust, saturation-based search procedure that is reminiscent of automated theorem provers.

We briefly review the basic algorithm of auto2 here. First, the statement to be proved is transformed into contradiction form $[A_1, \ldots, A_n] \implies \mathsf{False}$. The algorithm maintains a list of *items*, which as a first approximation can be thought of as propositions that follow from the assumptions $A_1, \ldots, A_n$. This list begins with the assumptions themselves. New items are created by *proof steps*, which are user-defined functions that match *one or two* existing items, and produce new items that logically follow from the matched ones. The new items are added to the main list in order of a computed priority, in the style of a best-first search. The search succeeds when a contradiction is found by one of the proof steps.

There are several modifications to this basic picture, including a box lattice for case analysis, and a rewrite table for E-matching. For these, we refer to [22] for details. In what follows, we discuss some of the more advanced features in auto2 that are not extensively covered in the original paper, but will be crucial to the current work.

Shadowing In addition to adding items, a proof step can also *shadow* one of the input items. Shadowing an item takes the item out of consideration by any proof step in the future. One application of shadowing is removing redundancies: if items A and B are equivalent, one of them can be shadowed. Shadowing can also be used to simulate changing the state of an item. For example, a proof step can, upon given input item A (and possibly another item B), shadow A and at the same time output a new item A', in effect changing A into A'.

Item types In the simple picture above, we thought of items as propositions. In reality, there are many types of items, and the user has the option of adding new item types. Each item type is given by a string, and each item contains a ty_str field that specifies its type. Each item also contains two data fields: tname (an Isabelle term) and prop (an Isabelle theorem). The interpretation of these two fields depend on the type of the item.

The two most common item types are `PROP` and `TERM`. The `PROP` items are plain propositions. Its tname is the statement of the proposition, and its prop is the theorem justifying the proposition from the initial assumptions. The `TERM` items represent terms that are relevant to the current proof. Its tname is the term itself, and its prop is the trivial true, as no formal justification is needed for considering a term. A more interesting item type, `DISJ`, will be discussed below. For our work on separation logic we will need to add three new item types.

Custom matching functions Most proof steps begin by matching its one or two input items to some pattern. The most basic kind of matching involves an E-matching (using the rewrite table) between a pattern p and the tname of a `PROP` item, returning for each match a proposition of form p. However, one can also specify custom matchers. This is especially important for items of type other than `PROP`.

Each custom matcher corresponds to an item type, and can only be used on items of that type. There are two kinds of custom matchers:

- A *prop-matcher* is a generalization of the kind of matching described above. It takes as input a pattern p and an item of the specified type (as well as the rewrite table), and attempts to derive propositions of form p from the prop field of the item. Each result it returns consists of an instantiation σ_i of arbitrary variables in p, and the derived proposition $p(\sigma_i)$.
- A *typed-matcher* is used for retrieving items according to its tname. It takes as input a pattern p and an item of the specified type, and matches p to the tname of the item in some custom manner. For each match, it returns the instantiations, as well as a theorem, usually the conjunction of equalities between parts of p and parts of the tname. However, the exact interpretation of both the input p and the output theorem can vary greatly depending on the item type.

Example: the `DISJ` item type We now illustrate the preceding concepts by discussing the item type `DISJ` (this is the combination of item types `DISJ` and `DISJ_ACTIVE` in [22]).

Items of type `DISJ` represent disjunctions, with or without arbitrary variables. Arbitrary variables (called schematic variables in Isabelle) are indicated by adding a "?" in front of its name. For example,

$$x < 2 \vee x > 3$$

is a proposition about a variable x already fixed in the current proof. On the other hand,

$$?x < 3 \vee ?x \geq (3 :: \alpha),$$

where $?x$ is a schematic variable, means $x < 3 \vee x \geq 3$ for all x of some fixed type α. Both propositions above can be represented by `DISJ` items.

The prop field of a `DISJ` item stores the disjunction theorem. The tname field is a tuple containing the list of disjuncts, preceded by two boolean flags is_active and match_prem. The meanings of the flags (to be realized by the proof steps and custom matchers below) are as follows: a `DISJ` item $P_1 \vee \cdots \vee P_n$, not containing schematic variables, will create a case analysis on P_1 only if is_active is on. A `DISJ` item with schematic variables never creates case analysis, but its is_active flag may be passed onto `DISJ` items derived from it. A `DISJ` item with match_prem flag on will be used by prop-matchers. They will not be used (for efficiency reasons) if the flag is off.

We now discuss some common proof steps working with `DISJ` items. To begin with, there are proof steps that create `DISJ` items from plain propositions. For example, from the proposition $P_1 \vee \cdots \vee P_n$ a proof step creates the `DISJ` item with

$$\mathsf{tname} = (\mathsf{true}, \mathsf{true}, P_1, \ldots, P_n), \quad \mathsf{prop} = P_1 \vee \cdots \vee P_n.$$

On the other hand, from the proposition $P_1 \longrightarrow P_2 \longrightarrow \cdots \longrightarrow P_n$ a proof step creates the `DISJ` item with

$$\mathsf{tname} = (\mathsf{false}, \mathsf{true}, \neg P_1, \ldots, \neg P_{n-1}, P_n), \quad \mathsf{prop} = \neg P_1 \vee \cdots \vee \neg P_{n-1} \vee P_n.$$

Note the is_active field is false in the second case, meaning this item will not create case analysis. In this way, disjunctions are treated differently from implications, even though they are logically equivalent. Another proof step can then determine whether case analysis will be performed on a disjunctive output, by outputing it as a disjunction or as an implication.

There are also proof steps deriving `DISJ` items from forall facts and existence goals. From $\forall x.\, P_1(x) \longrightarrow P_2(x)$ a proof step creates the item

$$\mathsf{tname} = (\mathsf{false}, \mathsf{true}, \neg P_1(?x), P_2(?x)), \quad \mathsf{prop} = \neg P_1(?x) \vee P_2(?x).$$

When schematic variables are present, `DISJ` items also cover the case where there is only one disjunct. For example, from $\neg(\exists x.\, P(x))$ a proof step creates the item

$$\mathsf{tname} = (\mathsf{false}, \mathsf{true}, \neg P(?x)), \quad \mathsf{prop} = \neg P(?x).$$

Next, we consider proof steps for reducing a disjunction. Given any `DISJ` item with $\mathsf{prop} = P_1 \vee \cdots \vee P_n$ without schematic variables, and a proposition $\neg P_1$, a proof step creates the `DISJ` item with $\mathsf{prop} = P_2 \vee \cdots \vee P_n$ (if $n = 2$, a plain proposition is created instead). It also shadows the original `DISJ` item. In effect, the proof step updates the original `DISJ` item by removing one of the disjuncts that is known to be false.

Given a `DISJ` item with $\mathsf{prop} = P_1(?x) \vee \cdots \vee P_n(?x)$, containing the schematic variable $?x$, and a proposition $\neg P_1(x)$, a proof step creates the `DISJ` item with

prop $= P_2(x) \vee \cdots \vee P_n(x)$. The original DISJ item is *not* shadowed, because it is not made redundant (there can be other instantiations of $?x$). In both cases, we keep the is_active field the same, but set match_prem to false: for efficiency reasons we do not match intermediaries in the reduction process in prop-matchers.

Finally, we consider the proof step for creating case analysis from DISJ items. Given a DISJ item $P_1 \vee \cdots P_n$, the proof step checks that the item contains no schematic variables, and that is_active is set to true. Then, it creates a new case with assumption P_1. If that case is resolved (that is, if $\neg P_1$ is derived), the reduction proof steps will reduce the item to $P_2 \vee \cdots \vee P_n$, and the next case P_2 will be created, and so on. In this way, we realize case analysis on the whole disjunction $P_1 \vee \cdots \vee P_n$.

We now consider custom matching functions for DISJ. The typed-matcher matches the input pattern to the list of disjuncts. For example, the pattern $?P \vee ?Q$ will match any DISJ item with at least two disjuncts (that is, all except those with schematic variables and a single disjunct). The prop-matcher is more complicated. First, it normalizes the input pattern to a disjunction with or without schematic variables. This normalized p is then matched against the prop field of the item, up to the AC-property of disjunction. If one of the disjuncts in p can be matched to anything (for example, $?P$ or $?P(?x)$, where $?x$ appears in the other disjuncts and $?P$ does not), care is taken to match that disjunct last, so it can be instantiated to whatever that remains after matching the other disjuncts.

In summary, the proof steps working with DISJ items form a module that creates, updates, and makes use of DISJ items. Together they implement heuristics such as the different treatment of $\neg A \vee B$ and $A \longrightarrow B$. Custom matching functions allow DISJ items to be used by other proof steps with little additional code. In what follows, we will build a similar module for separation logic around three new item types.

3 Description of the system

3.1 Example: ordered insertion preserves sortedness

Before delving into the details, we illustrate the style of reasoning we would like to emulate using an informal proof. The goal is to prove that ordered insertion (code given in Section 2.1) preserves sortedness. The Hoare triple is:

```
<os_list xs b ∗ ↑(sorted xs)>
os_insert x b
<λr. ∃A xs'. os_list xs' r ∗ ↑(sorted xs')>"
```

We assume that the following Hoare triple for the behavior of os_insert on the set of values is already proved:

```
<os_list xs b>
os_insert x b
<λr. ∃A xs'. os_list xs' r ∗ ↑(set xs' = {x} ∪ set xs)>
```

Moreover, we restrict to the most interesting case, where $b = \mathsf{Some}\ p$ and $x > \mathsf{val}\ v$.

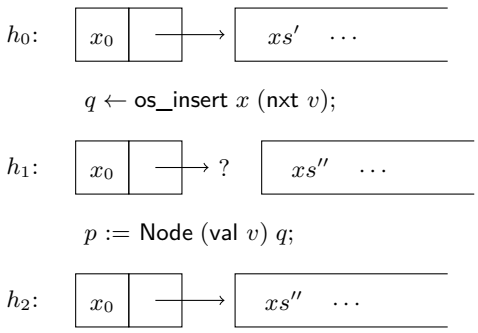

Fig. 1. Memory layout during ordered insertion

In the first stage of the proof, we map the memory layout of each heap encountered during the execution of the procedure, and name any data structures on the heaps (see Figure 1). The initial heap h_0 contains a linked list representing $xs = x_0 :: xs'$. After the recursive call to insert, the linked list for xs' is modified to that for xs'' in h_1. The head address of that list is possibly modified as well, so the nxt field of the node containing x_0 is not necessarily pointed there. This is remedied by the second line of the code. After which the heap h_2 contains a linked list for $x_0 :: xs''$. The head of this list is the return value of the procedure.

In the second stage of the proof, we derive facts about data structures on the heaps, starting with the pre-condition as assumptions, and the post-condition as a goal. A possible derivation is as follows:

1. $xs = x_0 :: xs'$ is sorted (pre-condition).
2. Suffice to show $x_0 :: xs''$ is sorted (post-condition).
3. (a): xs' is sorted, and (b): $x_0 \leq y$ for all $y \in \mathrm{set}\, xs'$ (1, def. of sortedness).
4. xs'' is sorted (3a and inductive hypothesis).
5. Suffice to show $x_0 \leq y$ for all $y \in \mathrm{set}\, xs''$ (2, 4, def. of sortedness).
6. $\mathrm{set}\, xs'' = \{x\} \cup \mathrm{set}\, xs'$ (Hoare triple for $\mathsf{os_insert}$ about set of values).
7. Suffice to show $x_0 \leq x$ and $x_0 \leq y$ for all $y \in \mathrm{set}\, xs'$ (5, 6, def. of union).
8. Suffice to show $x_0 \leq x$ (3b, 7).
9. Qed (8, $x > x_0$ since we are in the else branch).

This concludes the proof. Let us examine the salient characteristics of this proof, which will guide the design of our automation framework.

– Assertions about heaps are separated into two parts. The first part, which we call *spatial*, concerns the memory layout of the heap. The second part, which we call *pure*, concerns properties of data structures on the heap. Variables for the data structures (like xs) form the link between these two parts.

For example, the assertion "h contains a sorted linked list xs at address p" is divided into the two parts as follows:

$$h \vDash \underbrace{\mathsf{os_list}\ xs\ p}_{\text{spatial}} * \underbrace{\uparrow (\mathsf{sorted}\ xs)}_{\text{pure}}.$$

- In the first stage of the proof, we map the memory layout of each heap encountered during the execution of the procedure. This involves creating variables for the heaps and for each data structure that appear. In this stage we use the spatial part of the Hoare triples (but do need to verify the pure parts in the pre-conditions). Entailment theorems on assertions may also be used at this stage to bring assertions on the relevant heap into a form that matches the pre-condition.
- In the second stage, we derive facts about the data structures that appear. Here, the pre-condition serves as the initial assumption, and the post-condition serves as a goal. The pure parts of Hoare triples serve as rules of reasoning. Any other results unrelated to separation logic (for example properties of abstract data structures) may also be used at this stage.
- When the procedure is recursive, the inductive hypothesis serves as one of the Hoare triples that may be used in both stages of the proof.
- Multiple Hoare triples for the same procedure that share the same spatial part but have different pure parts may be used at the same time, to derive facts about the same data variables (here precision of certain assertions plays an important role, as we will make clear later). In the given example, we used the inductive hypothesis as well as the behavior on the set of values for os_insert.

3.2 Item types

In this section, we describe the three new item types defined in our automation framework. From now on, we will refer to partial heaps simply as heaps.

- `CODE_POS`: an item of type `CODE_POS` marks a position in the code reached by the prover. It is produced when all previous lines of code are known to run successfully. The tname field of a `CODE_POS` item is a pair (h_{i-1}, c_i), where h_{i-1} is the current (partial) heap, and c_i is the next line of code. Let $c' = c_i; \ldots; c_n$ be the complete remaining code. Then the prop field of the item contains the following fact: if the execution of c' on h_{i-1} is successful resulting in heap h' and return value r, then it suffices to prove $h' \vDash Q(r)$, where Q is the post-condition of the overall goal.
- `MATCH_POS`: an item of type `MATCH_POS` represents a successful match between a `CODE_POS` item and a Hoare triple $\langle P \rangle\, c\, \langle Q \rangle$, where c is the next line of code. Its tname is a tuple $(h_{i-1}, c_i, b_{\text{success}}, P, Q)$, where h_{i-1} and c_i come from the `CODE_POS` item. The meaning of the boolean flag b_{success} will be made clear below. The assertions P and Q are the pre/post-conditions of the Hoare triple. The prop field of the item is a conjunction of the prop field of the corresponding `CODE_POS` item and the Hoare triple.

– SUCCESS_RUN: an item of type SUCCESS_RUN indicates the successful execution of a line of code. Its tname is a tuple (h_{i-1}, c_i) where c_i is the relevant line of code, and h_{i-1} is the heap before the execution of c_i. Such an item is always derived from a Hoare triple $\langle P \rangle\, c_i\, \langle Q \rangle$. The prop field is the conjunction of three facts: that c_i carries h_{i-1} to h_i (the new heap), and the two assertions $h_{i-1} \vDash P$ and $h_i \vDash Q$.

3.3 Main proof steps

We now describe the main proof steps in our framework. Note all proof steps are invoked by auto2 automatically, so the user does not need to remember their names (although that may be helpful when reading the proof trace). All proof step names are also preceded by "sep." which we omit in what follows.

Throughout this description, we will follow the example given in Section 3.1, indicating the role each proof step plays in simulating the informal proof given there.

Initialization When the statement to be proved is a Hoare triple, the proof step init_pos creates an initial CODE_POS item. This proof step also creates a variable for the initial heap, processes the pre-condition, and stores a version of the goal in the prop field of the CODE_POS item.

For example, when given the goal

```
<os_list xs b ∗ ↑(sorted xs)>
os_insert x b
<λr. ∃A xs'. os_list xs' r ∗ ↑(sorted xs')>,
```

it creates variable h_0 for the initial heap, facts $h_0 \vDash$ os_list xs b and sorted xs, and the CODE_POS item with heap h_0 and c_1 being the entire code of os_insert. The prop field of the CODE_POS item contains the fact that if executing the entire code of os_insert x b on h_0 is successful with resulting heap h' and return value r, then it suffices to prove $h' \vDash \exists_A xs'.$ os_list xs' $r \ast \uparrow$ (sorted xs').

Rewrite to bind form Given a CODE_POS item of form (h_{i-1}, c_i) where c_i can be rewritten to a bind form $c_i'; c_i''$, the proof step rewrite_pos produces a new CODE_POS item (h_{i-1}, c_i'), where c_i is also rewritten to $c_i'; c_i''$ in the complete remaining code.

This proof step is mainly useful for handling if and case constructions in the code. For example, when analyzing the code for os_insert, we first split into cases where $b =$ None and $b =$ Some p. In the first case, the code is equivalent to os_prepend x None, and the case can be resolved using an existing Hoare triple for os_prepend. In the second case, the code is equivalent to $v \leftarrow\, !p; \ldots$. This proof step then creates a CODE_POS item with heap h_0 and $c_1 = !p$.

Likewise, after stepping through $v \leftarrow\, !p$ (which creates a variable v for the node at address p, and leaves the heap unchanged), case analysis will be performed on $x \leq$ val v. The first case is again resolved using a Hoare triple for

os_prepend and facts about sorted. In the second case, the code is equivalent to $q \leftarrow$ os_insert x (nxt v); .... This proof step then creates the CODE_POS item

$$(h_0, \text{os_insert } x \ (\text{nxt } v)). \tag{1}$$

Matching with Hoare triple We now discuss proof steps that match a CODE_POS item with a Hoare triple, returning a MATCH_POS item for each successful match.

For each previously proved Hoare triple $\langle P \rangle\, c\, \langle Q \rangle$, one can annotate the theorem to create a proof step that, when given a CODE_POS item (h_{i-1}, c_i), matches the pattern c to term c_i. For each successful match, it instantiates the matched variables in the Hoare triple, so the theorem becomes $\langle P' \rangle\, c_i\, \langle Q' \rangle$. Then it creates a MATCH_POS item with tname $(h_{i-1}, c_i, \text{false}, P', Q')$, and whose prop field is the conjunction of the prop for the original CODE_POS item, and the Hoare triple $\langle P' \rangle\, c_i\, \langle Q' \rangle$.

There are also two proof steps for matching with Hoare triples derived during the proof. The proof step match_hoare_prop matches CODE_POS items with Hoare triples without schematic variables. The proof step match_hoare_disj handles the case with schematic variables (where the Hoare triple is represented by a DISJ item).

In our example, focusing on the $b = \text{Some } p$ case. We get $xs \neq []$. By induction on xs, we get the inductive hypothesis (as a DISJ item):

```
<os_list (tl xs) ?b ∗ ↑(sorted (tl xs))>
os_insert x ?b
<λr. ∃A xs'. os_list xs' r ∗ ↑(sorted xs')>
```

From this and the CODE_POS item (1), match_hoare_disj will find a match with $?b :=$ nxt v, resulting in the MATCH_POS item

$$\begin{aligned}&(h_0, \text{os_insert } x \ (\text{nxt } v), \text{false}, \text{os_list } (\text{tl } xs) \ (\text{nxt } v) * \uparrow (\text{sorted } (\text{tl } xs)),\\ &\lambda r. \exists_A xs'. \text{os_list } xs' \ r * \uparrow (\text{sorted } xs'))\end{aligned} \tag{2}$$

In addition, the proof step created as the result of annotating the Hoare triple

```
<os_list xs b>
os_insert x b
<λr. ∃A xs'. os_list xs' r ∗ ↑(set xs' = {x} ∪ set xs)>
```

also matches (1), resulting in the MATCH_POS item:

$$\begin{aligned}&(h_0, \text{os_insert } x \ (\text{nxt } v), \text{false}, \text{os_list } ?xs \ (\text{nxt } v),\\ &\lambda r. \exists_A xs'. \text{os_list } xs' \ r * \uparrow (\text{set } xs' = x \cup \text{set } ?xs))\end{aligned} \tag{3}$$

Note in this case the arbitrary variable $?xs$ remains uninstantiated in the result.

First round matching with assertion A `MATCH_POS` item with $b_{\text{success}} = \mathsf{false}$, and whose pre-condition does not contain any pure assertions, can be used to derive the successful execution of the next line of code, by matching with an appropriate assertion about the heap. We call such a matching (for deriving successful execution) a *first round matching*.

Specifically, given a `MATCH_POS` item with tname $(h, c, \mathsf{false}, P, Q)$, and another item asserting the proposition $h \vDash P_0$, the match_assn1 proof step matches P_0 with the pattern $P * ?P'$ (here $?P'$ is an arbitrary variable, and P itself may contain arbitrary variables). This matching uses the associative-commutative property of separating conjunction, as well as known equalities in the rewrite table. For each match, the proof step produces a `SUCCESS_RUN` item with tname (h, c), a variable h' for the next heap, a variable for the return value (if any), and the `CODE_POS` item corresponding to the next line of code. The proof step also processes the post-condition, creating variables for each $\exists_A$ quantifier in the assertion, and splitting the assertion into spatial and pure parts. The results are stored in the prop field of the `SUCCESS_RUN` item, but not yet output as items themselves (the output of post-condition is left to the second round).

For example, from the pre-condition $h_0 \vDash \mathsf{os_list}\ xs\ b$ and $xs = x_0 :: xs'$, we get $h_0 \vDash p \mapsto_r \mathsf{Node}\ x_0\ n * \mathsf{os_list}\ xs'\ n$. Here $b = \mathsf{Some}\ p$, $v = \mathsf{Node}\ x_0\ n$, and $n = \mathsf{nxt}\ v$. The assertion matches the pattern $\mathsf{os_list}\ ?xs\ (\mathsf{nxt}\ v) * ?P'$ from the pre-condition of (3). Hence the line $\mathsf{os_insert}\ x\ (\mathsf{nxt}\ v)$ can be successfully executed. The proof step creates variables h_1 for the resulting heap and q for the return value. It derives the next `CODE_POS` item with tname $(h_1, p := \mathsf{Node}\ (\mathsf{val}\ v)\ q)$. The post-condition gives (note use of frame rule):

$$h_1 \vDash p \mapsto_r \mathsf{Node}\ x_0\ n * (\exists_A xs'.\ \mathsf{os_list}\ xs'\ q * \uparrow (\mathrm{set}\ xs' = \{x\} \cup \mathrm{set}\ xs))$$

This is processed as follows: the $\exists_A$ quantifier is moved to the front, variable xs' is created, and the assertion is split into spatial and pure parts, resulting in

$$h_1 \vDash p \mapsto_r \mathsf{Node}\ x_0\ n * \mathsf{os_list}\ xs'\ q \quad \text{and} \quad \mathrm{set}\ xs' = \{x\} \cup \mathrm{set}\ xs.$$

This is stored as part of the conjunction in the prop field of the `SUCCESS_RUN` item (the other parts of the conjunction are the assertion on h_0 used, and the statement that execution of $\mathsf{os_insert}\ x\ (\mathsf{nxt}\ v)$ carries h_0 to h_1).

If the end of code is reached, the new heap h' is the final heap. Instead of creating the next `CODE_POS` item, the proof step uses the prop stored in the current `CODE_POS` item to create a goal on h' corresponding to the post-condition of the theorem to be proved.

Reduction of pre-condition The use of `MATCH_POS` items whose pre-condition contains pure parts is more complicated. The pure parts need to be reduced before the item can be matched with an assertion.

The proof step match_assn_pure is used to reduce the pure parts one-by-one. Given a `MATCH_POS` item whose pre-condition is of the form $P * \uparrow (b)$ (up to rearrangement using the AC-property of $\cdot * \cdot$), and a proposition equivalent to b,

the proof step outputs a new MATCH_POS item with pre-condition P, keeping all other fields the same. The input MATCH_POS item is shadowed. This is similar to the reduction of DISJ items in the case without schematic variables.

For example, from item (2) and the proposition sorted (tl xs) (which is derived from sorted xs by a proof step unfolding the definition of sorted), the proof step match_assn_pure creates the MATCH_POS item

$$\begin{aligned}&(h_0, \mathsf{os_insert}\ x\ (\mathsf{nxt}\ v), \mathsf{false}, \mathsf{os_list}\ (\mathsf{tl}\ xs)\ (\mathsf{nxt}\ v),\\ &\lambda r.\ \exists_A xs'.\ \mathsf{os_list}\ xs'\ r * \uparrow (\mathsf{sorted}\ xs'))\end{aligned} \tag{4}$$

The above proof step only reduces when b contains no schematic variables. When there are schematic variables in b, it is usually more efficient to first instantiate those schematic variables by matching the spatial part with an assertion. We make the assumption that all schematic variables in the pure part also appear in the spatial part. Specifically, the match_assn_spatial proof step performs the following task: given MATCH_POS item $(h, c, b_{\text{success}}, P, Q)$, where P contains schematic variables, and an assertion $h \vDash P_0$, it matches P_0 to the pattern $P_S * ?P'$, where P_S is the spatial part of P. For each match, it creates the MATCH_POS item with the schematic variables instantiated. The pure parts in the resulting item can then be reduced by match_assn_pure later.

Conversion to second round Once a MATCH_POS item is used to derive a SUCCESS_RUN item that asserts its command is run successfully, we need to prevent other MATCH_POS items for the same command from doing the same, because otherwise multiple variables for the next heap will be created. The proof step convert_match_pos manages this, by matching a MATCH_POS item of form $(h, c, \mathsf{false}, P, Q)$ to a SUCCESS_RUN item of form (h, c), creating the MATCH_POS item $(h, c, \mathsf{true}, P, Q)$ and shadowing the original item, in effect toggling the b_{success} field of the input item.

In our example, the following two MATCH_POS items are created from items (3) and (4):

$$\begin{aligned}&(h_0, \mathsf{os_insert}\ x\ (\mathsf{nxt}\ v), \mathsf{true}, \mathsf{os_list}\ ?xs\ (\mathsf{nxt}\ v),\\ &\lambda r.\ \exists_A xs'.\ \mathsf{os_list}\ xs'\ r * \uparrow (\mathrm{set}\ xs' = x \cup \mathrm{set}\ ?xs))\end{aligned} \tag{5}$$

$$\begin{aligned}&(h_0, \mathsf{os_insert}\ x\ (\mathsf{nxt}\ v), \mathsf{true}, \mathsf{os_list}\ (\mathsf{tl}\ xs)\ (\mathsf{nxt}\ v),\\ &\lambda r.\ \exists_A xs'.\ \mathsf{os_list}\ xs'\ r * \uparrow (\mathsf{sorted}\ xs'))\end{aligned} \tag{6}$$

While this proof step performs no logical reasoning, it plays a crucial role in allowing multiple Hoare triples for the same procedure, ensuring that at most one of them will be used to derive that the command is run successfully and create variables for the next heap.

Second round matching with assertion Given a SUCCESS_RUN item (h, c), indicating that execution of c on heap h is successful, resulting in heap h', a second round of matching derives assertions on h' from post-conditions of Hoare

triples. This is managed by the proof step match_assn2. This proof step matches a MATCH_POS item of form $(h, c, \mathsf{true}, P, Q)$ with a SUCCESS_RUN item of form (h, c). Using the pre-condition on h, and the fact that c carries h to h', both stored in the prop field of the SUCCESS_RUN item, the proof step derives an assertion on h' corresponding to Q (if P does not match the pre-condition on h, the proof step returns no output). The post-condition is further processed as in the first round, except we also make sure that variables for data structures are shared, using the precision property of the relevant assertions.

In our example, from item (5), we get exactly the post-condition stored in the SUCCESS_RUN item:

$$h_1 \vDash p \mapsto_r \mathsf{Node}\ x_0\ n * \mathsf{os_list}\ xs'\ q \quad \text{and} \quad \mathrm{set}\, xs' = \{x\} \cup \mathrm{set}\, xs.$$

From item (6), we get

$$h_1 \vDash p \mapsto_r \mathsf{Node}\ x_0\ n * \mathsf{os_list}\ xs'\ q \quad \text{and} \quad \mathsf{sorted}\ xs'.$$

Note the sharing of variable xs'. This is possible because the assertion on h_1 is compared with what is stored in the SUCCESS_RUN item, and the precision property of os_list:

$$h \vDash \mathsf{os_list}\ xs_1\ q * P_1 \implies h \vDash \mathsf{os_list}\ xs_2\ q * P_2 \implies xs_1 = xs_2$$

is used to reconcile the variables. Here we see the crucial use of precision to avoid creating multiple variables for the same data structure.

Direct matching with assertion For some commands, only one Hoare triple is necessary, and that Hoare triple does not contain pure parts in the pre-condition. In such cases, the two-stage matching process discussed above is unnecessary. For example, the Hoare triple for the lookup command is:

$$\langle p \mapsto_r x \rangle\ !p\ \langle \lambda r.\ p \mapsto_r x * \uparrow (r = x) \rangle$$

One can annotate such a Hoare triple so it is used directly. In the example of lookup command, the resulting proof step matches a CODE_POS item of form $(h, !p)$, and a proposition of form $h \vDash p \mapsto_r x * P'$, and creates variable r for the return value, the proposition $r = x$, and the CODE_POS item for the next line of code (here no new heap variable is necessary because the lookup command is heap-preserving).

3.4 Other proof steps

Apart from the above proof steps for working with Hoare triples, there are other proof steps in our framework, mostly for working at the assertion level. We discuss two important examples here.

Forward entailment Given an entailment theorem $A \Longrightarrow_A B$, where A and B are assertions, we get

$$h \vDash A * P \implies h \vDash B * P$$

for any heap h and assertion P. One can annotate the entailment theorem to create a proof step that matches a proposition of form $h \vDash A * P$ and outputs the proposition $h \vDash B * P$. As usual, the matchings are up to the associative-commutative property of separating conjunction.

Backward entailment We can also apply an entailment theorem in the backward direction. In particular, let A_S and A_P be the spatial and pure parts of A, so the entailment theorem is of the form $A_S * A_P \Longrightarrow_A B$. Then one can write a proof step that matches two propositions $h \vDash A_S * P$ and $\neg h \vDash B * Q$, and creates the proposition $\neg((h \vDash A_S * Q) \wedge A_P)$. The meaning is as follows: suppose we want to show that part of h satisfies B, with the remaining part satisfying Q, and we already know that part of h satisfies A_S, then it suffices to show that the remaining part satisfies Q, and that the pure part A_P also holds.

3.5 Custom matching for assertions

One complexity we face when working with data structures such as linked lists and trees is that there are often many ways to write the spatial part of the assertion on the heap. The most common choices are whether to expand the inductive definition of linked lists and trees. For example, given

$$h \vDash \mathsf{os_list}\ xs\ b,$$

where xs is known to be nonempty, we may write $xs = x_0 :: xs'$, $b = \mathsf{Some}\ p$, and get

$$h \vDash p \mapsto_r \mathsf{Node}\ x_0\ n * \mathsf{os_list}\ xs'\ n,$$

where n is a new variable for the next pointer. While we can store both assertions on h (multiple assertions on the same heap are permitted in our framework), this could result in a proliferation of assertions, especially when we get to trees, and when there are multiple such data structures on the heap. Another option, which we take in the current implementation, is to store the expanded assertion only, shadowing the non-expanded one. This, however, requires strengthening the matching procedure so that a pattern in the non-expanded form can also be matched. We implemented this functionality as a custom prop-matcher on `PROP` items. The matcher uses, in addition to E-matching up to rewrite table, the inductive definition of data structures. Such inductive definitions (for example, for linked lists and binary search trees) can be registered when they are first stated.

In the example above, only the expanded assertion will be kept. However, the prop-matcher, when given the assertion and a pattern $h \vDash \mathsf{os_list}\ ?xs\ b$, will be able to return a result with $?xs := x_0 :: xs'$.

4 Case studies

In this section, we describe the case studies performed using our automation framework. Among the basic examples, we verified:

- Linked lists: reversal, insertion, delete, merge, and the higher-order operations map, filter, and fold.
- Binary search trees: insertion, deletion, and search.
- Arrays: reverse and quicksort (these do not really use separation logic, as everything is stored at a single address. They are mostly copied from an earlier version [22] without separation logic).

Below we describe the verification of two more advanced data structures: red-black trees (insertion and search only) and indexed priority queues (all operations).

In all of our examples, we use the automation to handle all reasoning related to separation logic, essentially reducing the problem to lemmas about abstract data structures. Some of these lemmas can also be proved automatically using auto2 (with the existing proof steps about arithmetic, sets, etc.). The more difficult lemmas require several intermediate steps to be given as hints. Once these lemmas are proved, the proof of Hoare triples is automatic except for possibly providing the induction scheme. The implementation, including the case studies, is available at `https://github.com/bzhan/auto2`.

4.1 Red-black trees

Red-black trees have been a popular test case for verifying data structures. We use an imperative implementation of insertion on red-black trees that is a translation from Okasaki's implementation for functional red-black trees [16]. This is different from the implementation commonly found in textbooks. In particular, it uses recursion instead of while loops, and has four cases in the balance procedure instead of five. However, it does modify the tree by updating pointers and allocates no extra memory, as expected for imperative programs.

We prove that insertion preserves the red-black invariant, sortedness, and behaves as expected on the set of keys, and on the mapping corresponding to the red-black tree. The large number of properties to be proved makes permitting multiple Hoare triples for the same procedure especially important.

For the proofs on sortedness, set of keys, and associated mapping, we follow the approach given in [15]. That is, we first prove that the balance and insertion functions behave as expected on the inorder traversal of the tree, from which the remaining properties follow easily.

We begin with the definition of abstract red-black trees:

```
datatype ('a, 'b) rbt =
  Leaf | Node (lsub: ('a, 'b) rbt) (cl: color) (key: 'a) (val: 'b) (rsub: ('a, 'b) rbt)
```

The red-black invariant is defined as a predicate on the red-black tree:

```
fun black_depth :: ('a, 'b) rbt ⇒ nat where
  black_depth Leaf = 0
| black_depth (Node l R k v r) = min (black_depth l) (black_depth r)
| black_depth (Node l B k v r) = min (black_depth l) (black_depth r) + 1

fun cl_inv :: ('a, 'b) rbt ⇒ bool where
  cl_inv Leaf = True
| cl_inv (Node l R k v r) = (cl_inv l ∧ cl_inv r ∧ cl l = B ∧ cl r = B)
| cl_inv (Node l B k v r) = (cl_inv l ∧ cl_inv r)

fun bd_inv :: ('a, 'b) rbt ⇒ bool where
  bd_inv Leaf = True
| bd_inv (Node l c k v r) = (bd_inv l ∧ bd_inv r ∧ black_depth l = black_depth r)

definition is_rbt :: ('a, 'b) rbt ⇒ bool where
  is_rbt t = (cl_inv t ∧ bd_inv t)
```

Likewise we define the set of keys, corresponding mapping, and sortedness on the red-black tree:

```
fun rbt_set :: "('a, 'b) pre_rbt ⇒ 'a set" where
  "rbt_set Leaf = {}"
| "rbt_set (Node l c k v r) = {k} ∪ rbt_set l ∪ rbt_set r"

fun rbt_sorted :: ('a::linorder, 'b) rbt ⇒ bool where
  rbt_sorted Leaf = True
| rbt_sorted (rbt.Node lt c k v rt) =
   ((∀x∈rbt_key_set lt. x < k) ∧ (∀x∈rbt_key_set rt. k < x) ∧ rbt_sorted lt ∧ rbt_sorted rt)

fun rbt_in_traverse_pairs :: "('a, 'b) pre_rbt ⇒ ('a × 'b) list" where
  "rbt_in_traverse_pairs Leaf = []"
| "rbt_in_traverse_pairs (Node l c k v r) =
   (rbt_in_traverse_pairs l) @ [(k, v)] @ (rbt_in_traverse_pairs r)"

definition rbt_map :: "('a, 'b) pre_rbt ⇒ ('a, 'b) map" where
  "rbt_map t = map_of_alist (rbt_in_traverse_pairs t)"
```

Now for definitions in separation logic. We define two levels of assertions. On the first level, only the structural conditions are asserted:

```
  btree Leaf p = ↑(p = None)
  btree (rbt.Node lt c k v rt) (Some p) =
     (∃A lp rp. p ↦r Node lp c k v rp ∗ btree lt lp ∗ btree rt rp)
  btree (rbt.Node lt c k v rt) None = false
```

The procedures balanceR, balance, rbt_ins, and rbt_insert are defined in sequence. We omit the code here. rbt_insert is the overall insertion procedure, while rbt_ins is an intermediate procedure defined by recursion. We prove that each of the functions behaves as expected with respect to the black-depth invariant, color invariant, and inorder traversal, ending with the following two Hoare triples:

```
theorem insert_is_rbt:
  <btree t p ∗ ↑(is_rbt t)>
  rbt_insert k v p
  <λr. ∃A t'. btree t' r ∗ ↑(is_rbt t')>

theorem insert_inorder_pairs:
  "<btree t p ∗ ↑(rbt_sorted t)>
   rbt_insert k v p
   <λr. ∃A t'. btree t' r ∗
    ↑(rbt_in_traverse_pairs t' = ordered_insert_pairs k v (rbt_in_traverse_pairs t))>"
```

The properties that rbt_insert preserves sortedness and behaves as expected on the associated mapping follow trivially from the latter Hoare triple. This verifies all of the important properties for insertion on red-black trees. However, for verifying algorithms that use red-black trees, it is better to define a second level of assertions, explicitly expressing the idea that a red-black tree represents a set (the set of its keys) or a mapping. The definitions are (omitting restrictions on type classes):

```
definition rbt_set :: 'a set ⇒ ('a, 'b) rbt_node ref option ⇒ assn where
  rbt_set S p = (∃A t. btree t p ∗ ↑(is_rbt t) ∗ ↑(rbt_sorted t) ∗ ↑(S = rbt_key_set t))

definition rbt_map :: ('a, 'b) map ⇒ ('a, 'b) rbt_node ref option ⇒ assn where
  rbt_map M p = (∃A t. btree t p ∗ ↑(is_rbt t) ∗ ↑(rbt_sorted t) ∗ ↑(M = rbt_tree_map t))"
```

With these definitions, we can automatically prove (invoking Hoare triples proved for assertions at the first level):

```
lemma rbt_insert_rule1:
  <rbt_set S b> rbt_insert k v b <rbt_set ({k} ∪ S)>

lemma rbt_insert_rule2:
  <rbt_map M b> rbt_insert k v b <rbt_map (M {k → v})>
```

In the end, we may remove proof steps corresponding to Hoare triples at the first level, so that auto2 will use exclusively the two Hoare triples above when reasoning about programs that use rbt_insert, hence achieving a high degree of modularity.

4.2 Indexed priority queues

In the second example, we verify the correctness of an indexed priority queue. This data structure is also verified in [9], as part of the verification of Dijkstra's shortest path algorithm. We follow the general method described there. Since this data structure may be less familiar, we will review the basic ideas.

A priority queue is a binary tree where each node contains a number (or an element of another type with comparisons). It satisfies the ordering condition that the value at each node must be less than or equal to the values at its child nodes. It is usually represented as an array, containing values of the tree reading from top to bottom and then left to right. We require the tree to be fully-filled

in the sense that the corresponding array has no gaps. The operations insert and delete-min can be performed in $O(\log n)$ time, using a sequence of swaps to restore the ordering condition.

In an indexed priority queue, each node contains a key-value pair, where the keys are distinct, and the values obey the usual ordering condition. There is, in addition, a mapping from keys to their indices in the array. This mapping allows one to immediately find the entry corresponding to any key, so that updating the value corresponding to a key can also be done in $O(\log n)$ time. Swaps in this case must be accompanied by a corresponding swap in the index. In our example, we restrict to keys that are natural numbers less than some fixed n, and store the mapping as an array of fixed length n.

An example of an indexed priority queue is shown in Figure 2.

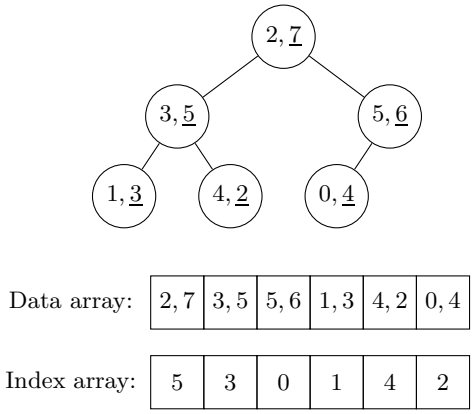

Fig. 2. Example of an indexed priority queue

We implement the indexed priority queue in a modular fashion, in terms of two smaller data structures. The first data structure, implementing the data array, is a *dynamic array*, which represents a resizeable array. It gives the illusion of having infinite capacity by reallocating memory whenever its capacity is exceeded. The second data structure is an array of fixed length n, interpreted as a mapping from integers $0, \ldots, n-1$. Each data structure is verified separately, resulting in Hoare triples for each of its operations. Some representative Hoare triples for dynamic array are as follows (the input x to dyn_array_new provides a default value for the type stored in the array).

```
theorem dyn_array_new_rule:
  <emp> dyn_array_new x <dyn_array []>

theorem push_array_rule:
  <dyn_array xs p> push_array x p <λr. dyn_array (xs @ [x]) r ∗ true>
```

```
theorem array_upd_rule:
  <dyn_array l p ∗ ↑(i < length l)>
  array_upd i x p
  <λ_. dyn_array (list_update l i x) p>
```

A selection of Hoare triples for the fixed length array, interpreted as a mapping from 0 to $n-1$, are as follows.

```
theorem amap_new_rule:
  <emp> amap_new k <λr. amap empty_map r ∗ ↑(alen r = k)>

theorem amap_lookup_rule:
  <amap m p> amap_lookup p k <λr. amap m p ∗ ↑(r = m⟨k⟩)>

theorem amap_update_rule:
  <amap m p ∗ ↑(i < alen p)> amap_update i x p <λ_. amap (m {i → x}) p>
```

The two data structures are put together to implement an indexed priority queue. As with red-black trees, we define two levels of assertions. At the first level, we have:

```
idx_pqueue xs (Indexed_PQueue pq idx) =
  (∃A m. dyn_array xs pq ∗ amap m idx ∗
   ↑(index_of_pqueue xs m) ∗ ↑(key_within_range xs (alen idx)))
```

This assertion states that the priority queue with data array pq and index array idx represents a list of key-value pairs xs, if the dynamic array pq represents xs, the index array idx represents a mapping m, such that m is actually the index of xs (index_of_pqueue predicate), and that each key in xs is within the limit of idx (key_within_range predicate).

At the second level, we have:

```
idx_pqueue_map M p =
  (∃A xs. idx_pqueue xs p ∗ ↑(is_heap xs) ∗ ↑(M = map_of_kv_list xs))
```

Here the is_heap predicate asserts that the list of key-value pairs xs obeys the ordering condition.

In the end, the Hoare triples proved at the second level are:

```
theorem delete_min_idx_pqueue_map:
  <idx_pqueue_map M p ∗ ↑(M ≠ empty_map)>
  delete_min_idx_pqueue p
  <λ(x, r). idx_pqueue_map (delete_map (key x) M) r ∗ ↑(∀k v. M⟨k⟩ = Some v → val x ≤ v)>

theorem insert_idx_pqueue_map:
  <idx_pqueue_map M p ∗ ↑(k < alen (index p)) ∗ ↑(M⟨k⟩ = None)>
  insert_idx_pqueue k v p
  <idx_pqueue_map (M {k → v}) ∗ true>

theorem update_idx_pqueue_map:
  <idx_pqueue_map M p ∗ ↑(M⟨k⟩ ≠ None)>
  update_idx_pqueue k v p
  <λ_. idx_pqueue_map (M {k → v}) p>
```

5 Related work

The current work is based on [22]. In that paper, we introduced the auto2 prover, and applied it to several case studies, including functional data structures and verification of imperative programs without using separation logic (or Hoare logic). The current work uses separation logic, and hence is able to simplify the proof process, and scale it up to more complex algorithms.

Automation in separation logic has been an area of extensive research. A lot of work has been done in the automatic theorem proving setting, where the focus is on finding ways to translate goals in separation logic into a form that can be handled by tools such as SMT solvers. Several of the resulting systems [14,17,18,19] are successful in verifying data structures such as linked lists and various kinds of binary search trees, including red-black trees. However, these systems are limited by restrictions to the expressiveness of the logic, either on the shape of the data structures (to linked lists and trees), or on the type of predicates that may be used. They are also for the most part restricted to first-order logic.

Closer to our approach are attempts to automate separation logic in a proof assistant such as Coq and Isabelle, by writing tactics that can perform a significant chunk of reasoning. Examples of these include [1,4,5,6,7,11,21]. The advantages of using proof assistants is that there is essentially no limit to the complexity of the program that can be verified. In addition, since the resulting proof is checked by a small kernel, one does not need to trust the implementation of the automation, placing a very high degree of confidence on the results. The drawback is that significant human guidance may be needed for the more complex verifications. It is also not designed to find counterexamples.

Of the existing work on verifying imperative implementations of algorithms, we mention especially the refinement framework [10]. It allows modular reasoning by first proving the correctness of an abstract version of the algorithm, then substituting data structures in the abstract version for their imperative implementations. This framework has been used to verify several graph algorithms [8,12,13]. We do not explicitly make use of the refinement framework in this paper, instead we aim for modularity in an ad-hoc manner.

The main difference between our work and the existing work in proof assistants is that our automation is implemented as an extension to auto2, rather than using the tactics framework. This means there is a significant difference in how search is conducted in our method. Generally speaking, there are two aspects that require search when using previously proved Hoare triples. First, the assertion on the heap may need to be rewritten into a form that matches the spatial part of the pre-condition of the Hoare triple. Such rewrites (or entailments) on assertions may come from definitions (such as the recursive definition of os_list), or proved lemmas (for example, that a linked list on $xs @ ys$ can be split into a list segment on xs and a linked list on ys). Second, the pure parts of the pre-condition need to be proved from the existing assumptions and derived facts. In automation implemented using tactics, such search is usually conducted by backtracking (if it is handled at all), essentially following a depth-first search

mechanism. This mechanism is susceptible to infinite loops or other dead ends, limiting the lemmas that can be automatically invoked. The best-first search mechanism that we use is much more robust, allowing essentially all previously proved entailments and Hoare triples to be applied automatically.

In our case studies, the amount of human guidance, as well as their complexity, is significantly lower than that for similar verifications using tactics-based automation, for example [21] for red-black trees and [9] for indexed priority queues. We believe that our more robust search mechanism, both for separation logic and for reasoning with abstract data structures, is most responsible for this improvement.

When verifying imperative programs, it is often convenient to use a functional specification as an intermediate layer. Appel gave in [2] an overview of the use of functional specification and its limitations. For the verifications in this paper, we intentionally avoid using functional specifications (they are used only for some small sub-procedures such as tree rotations). In addition to shortening the formalization, we wish to demonstrate the ability of our framework to handle the complexity of reasoning about separation logic and functional behaviors at the same time. Hence, our approach could be especially useful in situations where it is impossible or inconvenient to give functional specifications.

6 Conclusion

In this paper, we introduced a new framework for automation in separation logic, and applied it to verify the correctness of data structures including red-black trees and indexed priority queues. In these examples, the automation is able to take care of nearly all reasoning involving separation logic, leaving human guidance only for lemmas about properties of abstract data structures.

This work also gives an example of a complex extension to the auto2 prover. Auto2 is not initially designed for separation logic, and indeed is first applied to simple case studies in mathematics and functional data structures. The framework described in this paper demonstrates the potential of adding domain-specific heuristics to auto2.

There are many ways in which the current work can be extended. In particular, user-experience can be improved by providing ways to visualize a proof (either failed or completed), for example by indicating which lines of the procedure have been shown to be successful, and what are the known assertions about each heap and data variable. This can be viewed as a part of the task for improving proof visualization in auto2 in general. Given the closeness of our approach to human reasoning, we are confident that such visualization can be very helpful to the user in identifying problems with a proof.

Specifically for separation logic, there are several ways by which we can expand the range of algorithms that can be easily verified. We would like to support annotated while loops. For reasoning with arrays, it is often helpful for partial heaps to contain only part of an array. We would also like to add frameworks for verifying the time and space-complexity of algorithms. More challenging would

be defining new concepts, and heuristics for working with them, in order to verify algorithms involving more complex memory sharing patterns.

Acknowledgements. The author thanks Adam Chlipala and Peter Lammich for discussions during this project. This work is done while the author is supported by NSF Award No. 1400713.

References

1. Appel, A.: Tactics for separation logic, 2006. http://www.cs.princeton.edu/~ appel/papers/septacs.pdf.
2. Appel, A.: Modular Verification for Computer Security. In: 2016 IEEE 29th Computer Security Foundations Symposium (CSF), Lisbon, 2016, pp. 1–8.
3. Bulwahn, L., Krauss, A., Haftmann, F., Erkök, L., Matthews, J.: Imperative Functional Programming with Isabelle/HOL. In Theorem Proving in Higher Order Logics, LNCS 5170, pp. 134–149, 2008.
4. Cao, J., Fu, M., Feng, X.: Practical Tactics for Verifying C Programs in Coq. In X. Leroy and A. Tiu (Eds.): CPP 2015, pp. 97–108, 2015.
5. Charguéraud, A.: Characteristic formulae for the verification of imperative programs. In: ICFP, pp. 418–430. ACM (2011)
6. Chlipala, A.: Mostly-automated verification of low-level programs in computational separation logic. In PLDI'11, pp 234–245, 2011.
7. Chlipala, A., Malecha, G., Morrisett, G., Shinnar, A., Wisnesky, R.: Effective Interactive Proofs for Higher-Order Imperative Programs. In Proceedings of the 14th ACM SIGPLAN International Conference on Functional Programming (ICFP'09), pp. 79–90. August 2009.
8. Lammich, P.: Verified efficient implementation of Gabow's strongly connected component algorithm. In: Klein, G., Gamboa, R. (eds.) ITP 2014. LNCS, vol. 8558, pp. 325–340. Springer, Heidelberg (2014)
9. Lammich, P.: Refinement based verification of imperative data structures. In J. Avigad and A. Chlipala (Eds.): CPP 2016, pp. 27–36, 2016.
10. Lammich, P.: Refinement to Imperative/HOL. In C. Urban and X. Zhang (Eds.): ITP 2015, LNCS 9236, pp. 253–269, 2015.
11. Lammich, P., Meis, R.: A separation logic framework for imperative HOL. In: Archive of Formal Proofs, Nov 2012. http://afp.sf.net/entries/Separation_Logic_Imperative_HOL.shtml. Formal proof development.
12. Lammich, P. Sefidgar, S.R.: Formalizing the Edmonds-Karp Algorithm. In: J.C. Blanchette and S. Merz (Eds.): ITP 2016, LNCS 9807, pp. 219–234, 2016.
13. Lammich, P., Tuerk, T.: Applying data refinement for monadic programs to Hopcroft's algorithm. In: Beringer, L., Felty, A. (eds.) ITP 2012. LNCS, vol. 7406, pp. 166–182. Springer, Heidelberg (2012)
14. Le, Q.L., Sun, J., Chin, W. Satisfiability Modulo Heap-Based Programs. In S. Chaudhuri and A. Farzan (Eds.): CAV 2016, pp. 382–404, 2016.
15. Nipkow, T.: Automatic Functional Correctness Proofs for Functional Search Trees. In: J.C. Blanchette and S. Merz (Eds.): ITP 2016, LNCS 9807, pp. 307–322, 2016.
16. Okasaki, Chris.: Red-black trees in a functional setting. J. functional programming, 9(4), 471–477. (1999)

17. Pek, E., Qiu, X., Madhusudan, P.: Natural Proofs for Data Structure Manipulation in C using Separation Logic. PLDI 2014, pp. 440–451. (2014)
18. Piskac, R., Wies, T., Zufferey, D.: Automating Separation Logic with Trees and Data. In A. Biere and R. Bloem (Eds.): CAV 2014, pp. 711–728, 2014.
19. Qiu, X., Garg, P., Ştefănescu, A., Madhusudan, P.: Natural proofs for structure, data, and separation. PLDI 2013, pp. 231–242. (2013)
20. Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: Proceedings of Logic in Computer Science (LICS), pp. 55–74. IEEE (2002)
21. Tuerk, T.: A separation logic framework for HOL. Technical report UCAM-CL-TR-799, University of Cambridge, Computer Laboratory, June 2011
22. Zhan, B.: AUTO2: a saturation-based heuristic prover for higher-order logic. In: J.C. Blanchette and S. Merz (Eds.): ITP 2016, LNCS 9807, pp. 441–456, 2016.